

# IXPUG-ISC16 Workshop

OpenMP  
Affinity... ..on  
KNL

**PRESENTED BY:**

Kent Milfeld

[milfeld@tacc.utexas.edu](mailto:milfeld@tacc.utexas.edu)

# Affinity

## Why do we need it?

- Some Apps run better with **fewer threads** than HW-threads.  
Execution with only 1, 2, or 3 threads per core.  
Execution with only a few threads per tile.  
Execution with fewer cores per quadrant (e.g. avoid #6)  
Allow threads to float on cores.

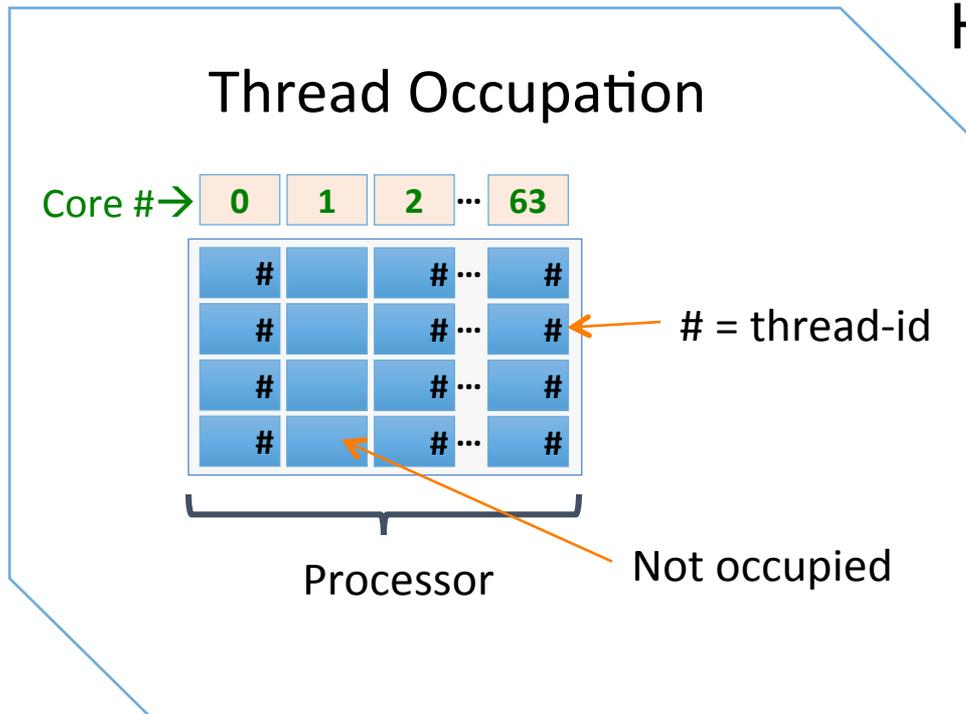
## What are the **OpenMP** tools?

- **distribution policy** (**close/spread** ...)
- **syntax for expressing** HW-threads (**places**) to run on

# Affinity (thread-id assignment)

In a parallel region threads are assigned to HW-threads.

We will now show occupation: assigned thread-ids on the HW-thread “grid”.



Layout is for a 64-core system

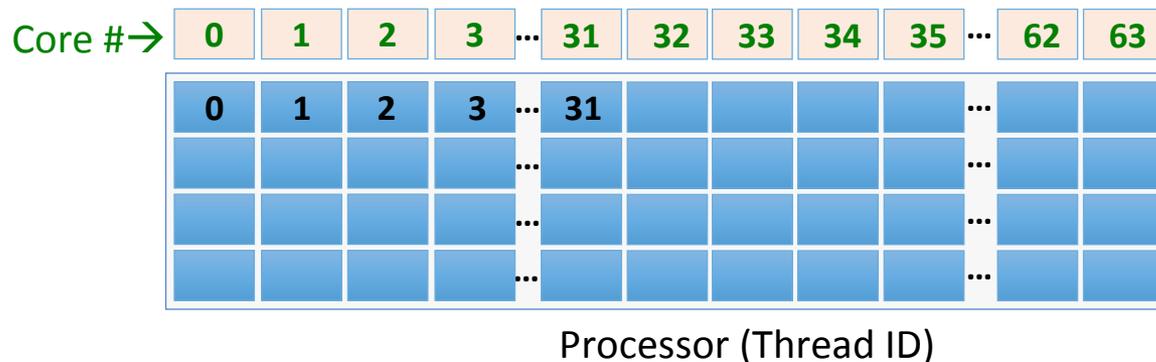
# Affinity (distribution)

Distributions (fewer threads than the total 64 cores):

**close** (keeps threads together)

`OMP_NUM_THREADS=32`

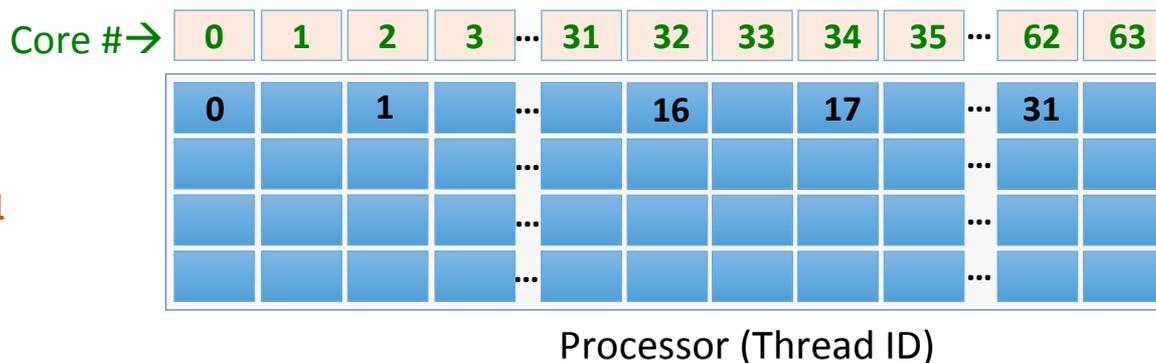
`OMP_PROC_BIND=close`



**spread** (spreads threads out across cores)

`OMP_NUM_THREADS=32`

`OMP_PROC_BIND=spread`



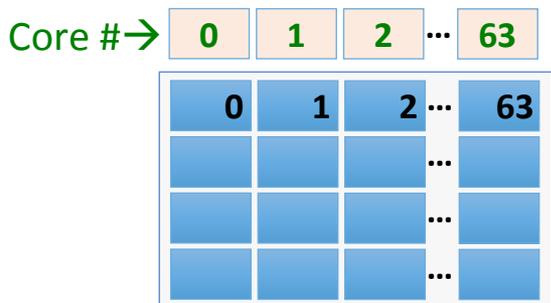
# Affinity (distribution)

Distributions (greater number of threads than the total 64 cores):  
Occupying 1, 2, 3 and 4 OpenMP threads per core:

```
export OMP_PROC_BIND=spread
```

1 thread/core

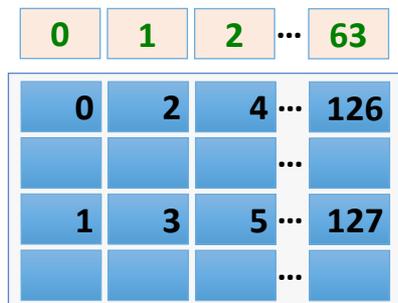
OMP\_NUM\_THREADS  
64



Processor (Thread ID)

2 threads/core

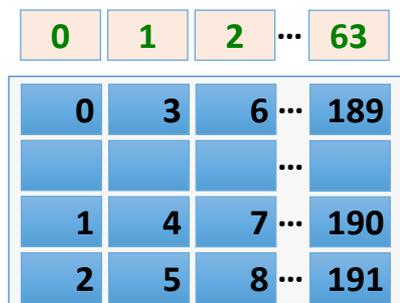
OMP\_NUM\_THREADS  
128



Processor (Thread ID)

3 threads/core

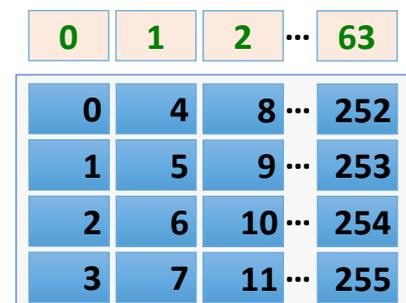
OMP\_NUM\_THREADS  
192



Processor (Thread ID)

4 threads/core

OMP\_NUM\_THREADS  
256



Processor (Thread ID)

Sequential threads are on the same core.

Try this affinity setting when increasing the threads per core.

# Affinity (distribution)

Distributions (greater number of threads than the total 64 cores):  
Occupying 1, 2, 3 and 4 of the quadrants:

```
export OMP_PROC_BIND=close
```

Quadrant 1

Quadrants 1-2

Quadrants 1-3

Quadrants 1-4

OMP\_NUM\_THREADS  
64

Core #→ 0 1 ... 15 ... 63

0	4	...	60	...	
1	5	...	61	...	
2	6	...	62	...	
3	7	...	63	...	

Processor (Thread ID)

OMP\_NUM\_THREADS  
128

0 1 ... 31 ... 63

0	4	...	124	...	
1	5	...	125	...	
2	6	...	126	...	
3	7	...	127	...	

Processor (Thread ID)

OMP\_NUM\_THREADS  
192

0 1 ... 47 ... 63

0	4	...	188	...	
1	5	...	189	...	
2	6	...	190	...	
3	7	...	191	...	

Processor (Thread ID)

OMP\_NUM\_THREADS  
256

0 1 ... 63

0	4	...	252
1	5	...	253
2	6	...	254
3	7	...	255

Processor (Thread ID)

# Affinity ( places)

- A PLACE is a list of HW-threads.

Comma separated list

e.g. {0},{1},{2},{3},{4}

- Basic place syntax: *{ lower-bound } : length : stride*

{ } : length

e.g. {0}:5 = {0},{1},{2},{3},{4}

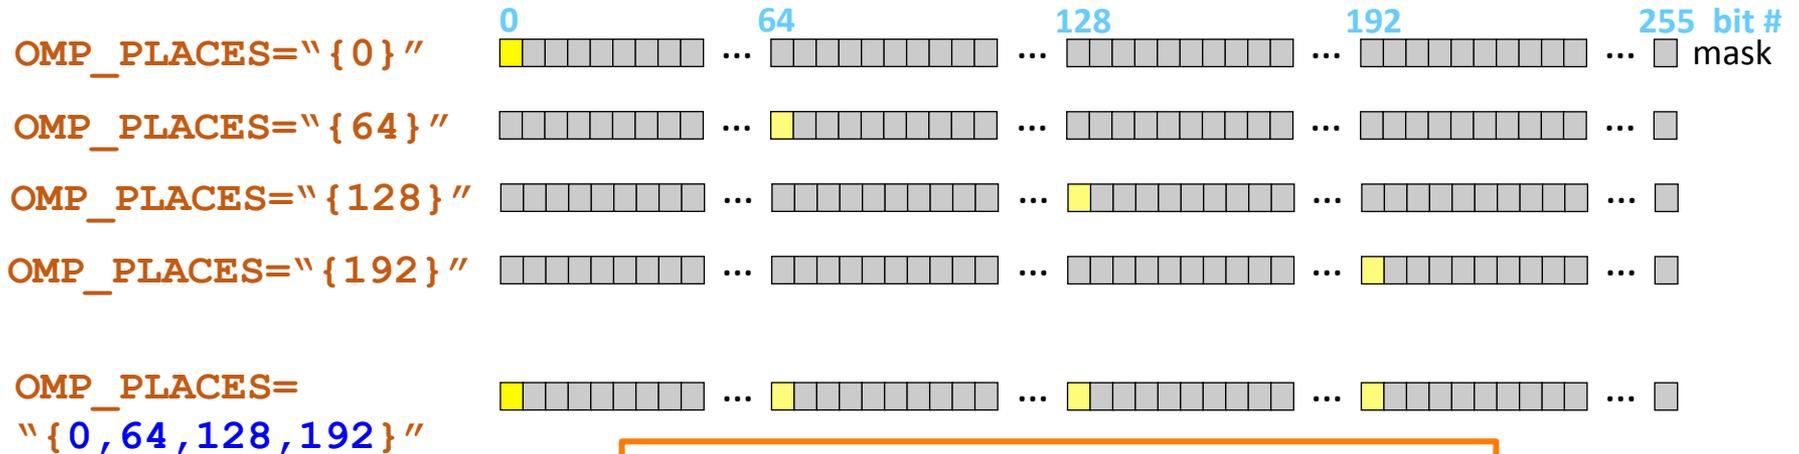
{ } : length : stride

e.g. {0}:3:2 = {0},{2},{4}

e.g. `export OMP_NUM_THREADS=128 OMP_PLACES="{0}:128"`

# Affinity (HW-thread maps, floating threads)

A 256-bit mask specifies where a thread may execute.



How can I get this mask (for the 1st core)?

A list within a place is a mask (where a thread can “float”).

# Affinity (list inside a place)

- A list inside a place creates a mask:

Comma separated list

e.g. { 0, 1, 2, 3, 4, 5 }

- Basic list syntax:

{ *lower-bound* : *length* : *stride* }

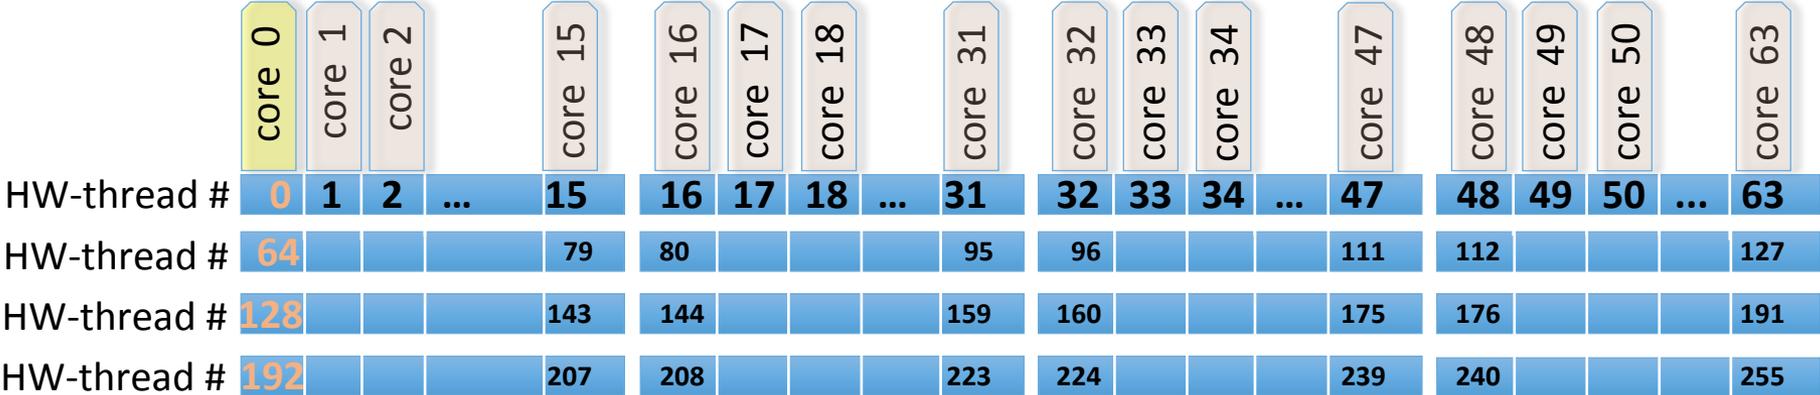
{ lb : length }

e.g. {0:4} = {0, 1, 2, 3}

{ lb : length : stride }

e.g. {0:3:2} = {0, 2, 4}

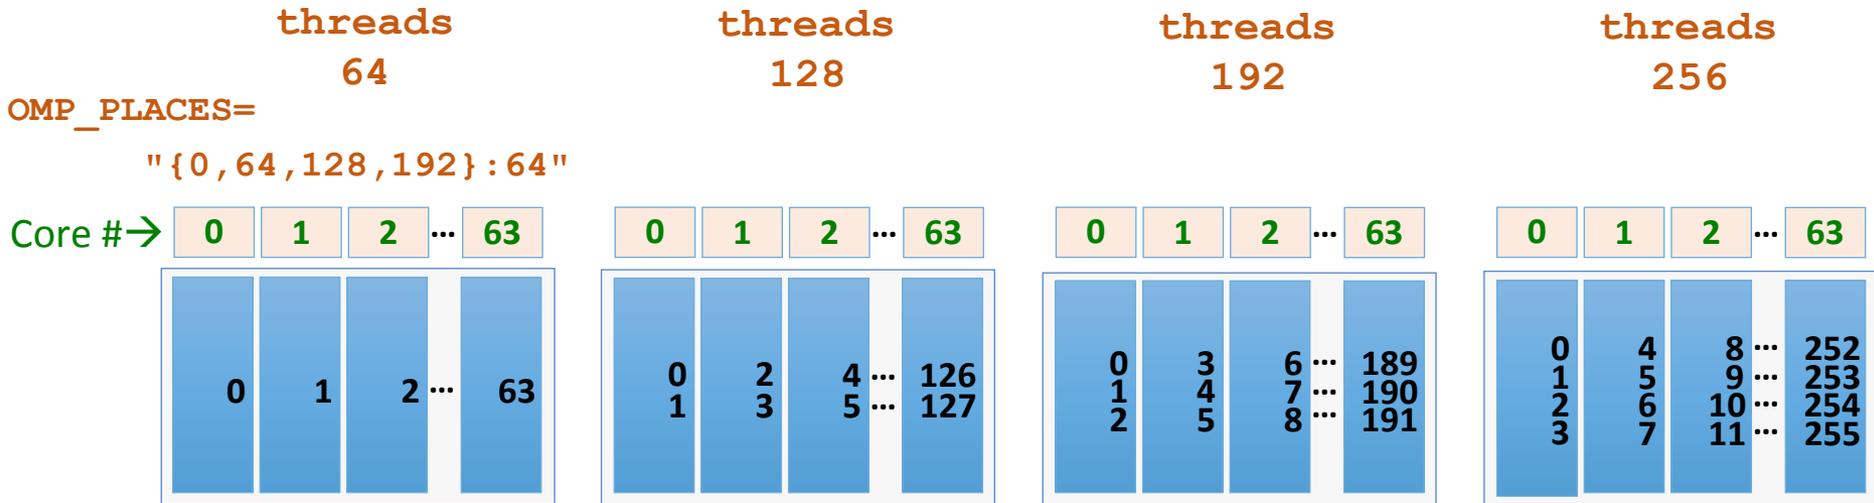
# Hardware Threads



Layout is for a 64-core system

# Affinity (cores)

List in a place allows a thread to float on any HW-thread (in a core here).



# Summary

- When using all KNL HW-threads, just **set the number of threads needed**– that may work fine.
- Try using **OMP\_PROC\_BIND=spread** for sequential thread ids on a core (for 1, 2, 3 and 4 threads per core).
- Use **OMP\_PLACES** for complex layouts of threads  
**OMP\_PLACES={*lower\_bound\**}:length:stride**  
syntax is a compact form for building HW-thread lists.

\* {lb} can be expanded as {lb:length:stride}

# The End

See me after the workshop for questions.

# Affinity (core summary)

- Affinity settings to assign 1, 2, 3 and 4 threads on a core.
- Threads in a core can execute on any HW-thread.
- Thread numbering is sequential on core.

OMP_NUM_THREADS	OMP_PLACES	Cores 0-15
64	{0,64,128,192}:64	1 thread/core
128	{0,64,128,192}:64	2 threads/core
192	{0,64,128,192}:64	3 thread/core
256	{0,64,128,192}:64	4 threads/core

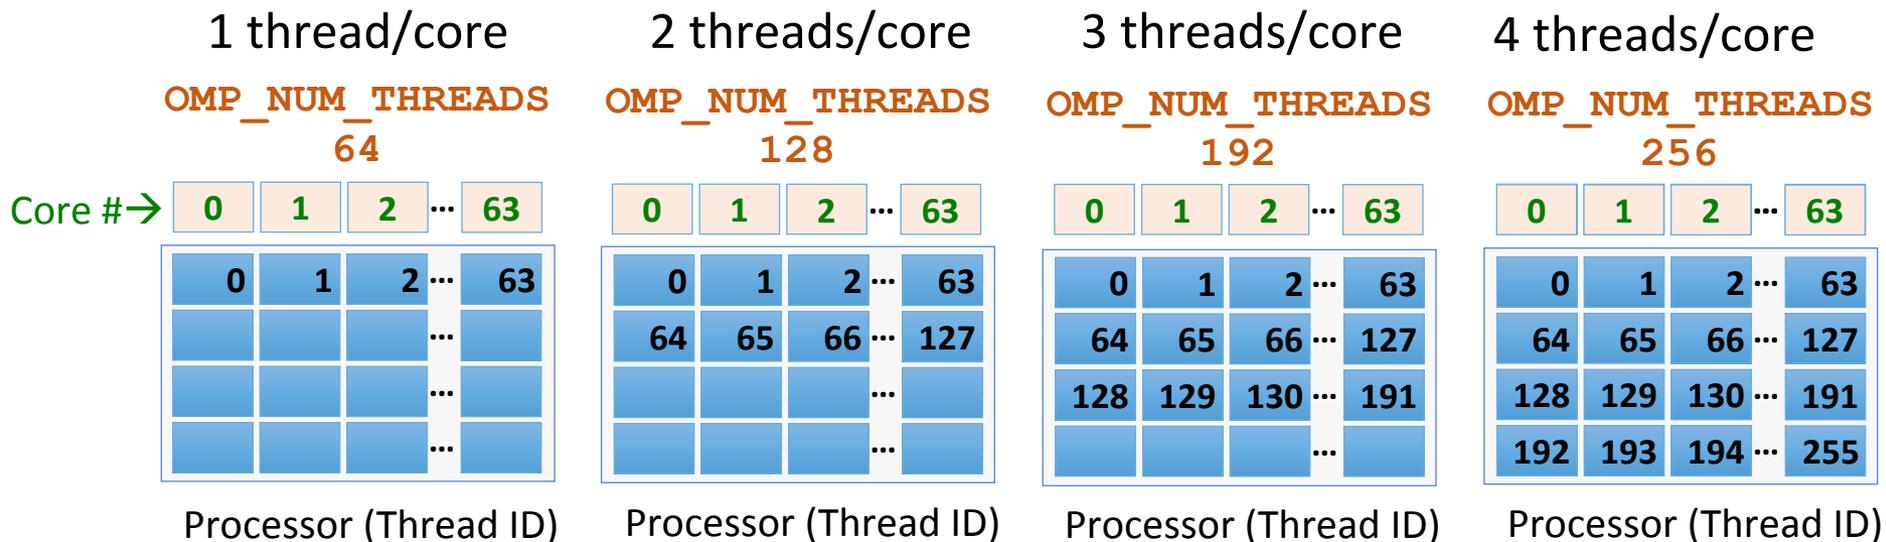
OMP\_PROC\_BIND does not need to be set.

e.g. `export OMP_NUM_THREADS=128 OMP_PLACES="{0,64,128,192}:64"`



# Default Affinity (distribution)

Occupying 1, 2, 3 and 4 HW threads per core:  
(Default: thread ID # → HW-thread id)



This is OK for some applications, but:

- 1.) Putting sequential threads on the same core may be more **cache** friendly,
- 2.) Allowing core threads to “float” in the core may be more **balance** friendly.

# Affinity (cores)

List in a place allows a thread to float on any HW-thread (in a core here).

